arXiv:1704.06676v1 [cs.AI] 21 Apr 2017

Multi-Objective Deep Q-Learning with Subsumption Architecture

Tomasz Tajmajer

Faculty of Mathematics, Informatics, and Mechanics,
University of Warsaw, Warsaw, Poland,
t.tajmajer@mimuw.edu.pl

Abstract In this work we present a method for using Deep Q-Networks (DQNs) in multi-objective tasks. Deep Q-Networks provide remarkable performance in single objective tasks learning from high-level visual perception. However, in many scenarios (e.g in robotics), the agent needs to pursue multiple objectives simultaneously. We propose an architecture in which separate DQNs are used to control the agent's behaviour with respect to particular objectives. In this architecture we use signal suppression, known from the (Brooks) subsumption architecture, to combine outputs of several DQNs into a single action. Our architecture enables the decomposition of the agent's behaviour into controllable and replaceable sub-behaviours learned by distinct modules. To evaluate our solution we used a game-like simulator in which an agent - provided with high-level visual input - pursues multiple objectives in a 2D world. Our solution provides benefits of modularity, while its performance is comparable to the monolithic approach.

Keywords: Deep Q-learning, Deep Reinforcement Learning, Subsumption Architecture, Neural Networks

1 Introduction

Many recent works on Reinforcement Learning focus on single-objective methods such as Deep Q-learning [1,2]. As those methods provide great performance in task such as playing video games, many real-life problems require satisfying multiple objectives simultaneously. In single objective reinforcement learning (SORL) the agent receives a single reward each time it performs an action. In multi-objective reinforcement learning (MORL) the agent receives multiple rewards - one for each objective. In particular, agents dealing with physical environment such as robots, need to pursue multiple, often conflicting objectives.

To have a graspable example, lets consider an autonomous cleaning robot, which is able to clean floors without colliding with any objects and return to charging station to recharge batteries. The observable aggregated behaviour of the robot may be decomposed into three sub-behaviours: collision avoidance (ca), floor cleaning (fc) and recharging (rg). We may thus describe the objectives of the robot for each identified sub-behaviour, to define the problem in a multi-objective manner, or we can aggregate the sub-behaviours and define a single

objective. In the former case, the robot-agent will receive a set of three rewards ($[r_{ca}, r_{fc}, r_{rg}]$) after each action. If the robot-agent collide with a wall, it may receive a negative reward related to collision avoidance (r_{ca}), but the rewards related to floor cleaning and recharging should not be influenced by this event. However, in the latter case, when we have a single objective, the robot-agent will receive only one reward value ($[r]$) related to any of the three sub-behaviours. Now, in case of collision, the robot-agent will also receive a negative reward, however a similar and indistinguishable negative reward may be provided if the agent e.g. deplete its batteries.

In single objective scenarios, we may find an optimal policy for which the sum of rewards collected by the agent is the highest possible. Methods such as Q-learning should converge to optimal policies [3]. However in case of multi-objective problems, there may be many such optimal policies depending on the trade-offs between satisfying particular objectives [4].

Autonomous agents, such as our example cleaning robot, are not really independent - they usually have a purpose defined by another agent: human. This aspect is often neglected in the literature, but is significant when considering practical applications of intelligent agents in robotics and automation. Our cleaning robot may follow a policy for which collision avoidance has greater importance than floor cleaning - in such case the robot should focus on avoiding collisions even at the cost of worse performance at floor cleaning. It is however for the user of such robot to decide, what should be the proportion between carefulness and cleanliness. The user may even want to fully disable some functions (behaviours) of the robot.

We see that when considering practical applications it is desired to have a multi-objective reinforcement learning method with the following features available post-learning: f.1) ability to select the sub-set of pursued objectives and f.2) ability to change the impact of particular objectives on the overall policy of the agent. Moreover, from the perspective of a manufacturer, it is also desired to have f.3) the ability to transfer learned behaviours between agents. As we will show later, the method presented in this paper posses all those features.

Multi-objective problems may be approached using *single-policy* or *multi-policy* methods. The simplest single-policy method uses a *scalarization function* [5], which converts multiple objectives into a single objective. Scalarization methods utilize a weight matrix to obtain a single score from multiple action-value functions. Some techniques assign linear priorities to objectives [6]. This allows to obtain a single optimal policy with respect to objectives ordered by those priorities.

In contrast to single-policy methods, multi-policy MORL methods are used for find a set of policies. Their aim is to approximate the Pareto front of policies [4]. In multi-policy methods, the preference of objectives does not need to be set a priori as a Pareto optimal policy for any preference may be obtained at runtime [7].

A natural approach in MORL is to use separate learning modules for each objective [8]. Modularity allows to decompose the problem into components that

are to some extent independent [9]; modularity may be required for providing features desired in practical applications that were listed earlier. Some works deal with transforming complex single-objective problems to many simpler objectives [10]. Such methods may be used to benefit from modular approach while solving single-objective problems.

Although Deep Q-Networks gained much attention in recent years, not many works consider the use of DQNs in multi-objective problems. Very recently authors of [11] proposed a multi-policy learning framework that utilizes Deep Q-Networks.

Learning behaviours in embodied agents, such as robots, is a problem well fitted for reinforcement learning methods. In *embodied artificial intelligence*, the idea of *parallel, loosely coupled processes* [12] is proposed as a principle for designing embodied agents. It states, that the control logic for embodied agents should consists of many independent components dedicated for particular aspects of the agent's behaviour. The aggregated behaviour of an agent emerges from cooperation or competence among those components. The *subsumption architecture* [13] is a very successful realization of the principle of parallel, loosely coupled processes. It is used in robotic designs from cleaning robots to mars rovers.

In the subsumption architecture, behaviours are divided into levels, implemented by distinct modules, forming a hierarchy of control. Modules of higher level may *subsume* modules of lower level, i.e. the output of lower level modules may be replaced by the output of higher level modules. Originally, modules were designed and implemented manually, and it was the role of the designer to decompose the expected behaviour of an agent into particular sub-behaviours. However, several methods for utilizing RL for behaviour learning were proposed [14,15]. In this work we will present a Deep Q-learning method for learning sub-behaviours by modules in the subsumption architecture.

First we will present how multiple DQNs may be used in parallel to control a single agent. Next, we will show how suppression connections may be implemented and how the control over agent may be shared by multiple DQNs. Finally, we will describe how to train multiple DQNs at once to promote modularity.

In the experimental section, we will present a simple 2D game - a virtual environment including an autonomous agent that has a local (situated) sensory inputs and may pursue different objectives. The game has different variants. We will use this game and its variants to evaluate and compare standard monolithic DQNs and our solution.

2 Background

2.1 Single Objective Reinforcement Learning

In the single-objective reinforcement learning an agent interacts with the environment by perceiving the state $s_t \in S$ and performing an action $a_t \in A$ for each step t. The actions are chosen by the agent according to some policy π. After performing an action, the agent receives a reward r_t. Then the agent observes

the next state s_{t+1} and the process repeats. The goal of the agent is to maximize the expected discounted reward $R_t = \sum_{k=0}^{\infty} \gamma^k r_{t+k}$, where $\gamma \in [0, 1]$ is the discount factor.

In Q-learning actions are selected based on $Q(s, a)$, which represents the expected discounted reward for performing action a in state s. For given state s, $a_t = \arg\max_a Q(s, a)$ is the optimal action. Deep Q-learning utilizes Deep Neural Networks for approximating Q(s,a) values, thus enabling this method to be used in many real-world applications. Deep Q-Networks [2] may be used used with high-level visual inputs such as those provided by video games.

2.2 Multi-Objective Reinforcement Learning

We may consider a more complex reinforcement learning scenario in which multiple objectives are pursued by the agent. Let O be the set of objectives of an agent. We may assign a priority to each objective $o \in O$ to form an linear sequence $[o_1, o_2, ..., o_n]$ such that o_k will have lower priority than o_j when $k < j$, where n is the number of objectives.

The agent, instead of a single reward, receives a vector of rewards at each time-step t with respect to each objective o_i, i.e: $\boldsymbol{r}_t = [r_{1,t}, r_{2,t}, \dots\ r_{n,t}]$, where $r_{i,t}$ corresponds to objective o_i. For each objective o_i and step t we may define the discounted return as:

$$R_{i,t} = \sum_{k=0}^{\infty} \gamma^k r_{i,t+k} \tag{1}$$

Moreover, for each objective o_i there is a Q-function $Q_i(s, a)$ that represents the expected discounted return $R_{i,t}$, i.e: $Q_i(s, a) = \mathbb{E}\left[R_{i,t} \,|\, s_t = s, a_t = a\right]$. We may define a vector of Q-functions, which includes $Q(s, a)$ for each objective o_i:

$$\boldsymbol{Q}(s, a) = [Q_1(s, a), Q_2(s, a), ..., Q_n(s, a)] \tag{2}$$

The function $Q_i(s, a)$ may be used by the agent to determine the optimal action with respect to objective o_i at time-step t, given state s_t:

$$a_{i,t} = \arg\max_a Q_i(s_t, a) \tag{3}$$

The vector $\boldsymbol{a}_t = [a_{1,t}, a_{2,t}, ..., a_{n,t}]$ consists of actions optimal with respect to particular objectives at given time-step t. At each step, the agent may perform only a single action, so a method of reducing $\boldsymbol{a}_t$ to a single action is required.

The most common method for selecting a single action is scalarization [5], which uses a weight vector $\boldsymbol{w}$ to retrieve a single scalar from vector $\boldsymbol{a}_t$. In the next section, we will show how a mechanism inspired by the subsumption architecture may be used instead.

3 Subsumption and Deep Q-Networks

We have considered an agent that have multiple objectives, receives rewards with respect to those objectives and has a separate Q-function for each objective. In

this section we will describe a) how actions obtained from different Q-function may be merged together, b) how an agent may switch between pursued objectives and c) how all Q-functions in the Q-vector may trained simultaneously.

We will refer to our method as to Multiple Deep Q-Network with Subsumption (mDQNS).

3.1 Combining multiple DQNs

In case of multi-objective agent, we may use a separate DQN as an approximator for each $Q_i(s,a)$ in the $\boldsymbol{Q}(s,a)$ vector. Such agent would be controlled by multiple Deep Q-Networks working in parallel. Each DQN provides a list of q-values and we want to use q-values from all DQNs to select a single action a that will be performed by the agent. Let us define a vector $\boldsymbol{q}_i$ that consist of q-values provided by $Q_i(s,a)$ for each possible action $a \in A$ and a single objective o_i, i.e.:

$$\boldsymbol{q}_i = [Q_i(s,a_0), Q_i(s,a_1), ..., Q_i(s,a_j)] \tag{4}$$

In the single-objective case the optimal action a would be equal to a_j for such j that $\boldsymbol{q}_{i,j} = \max \boldsymbol{q}_i$. Here we have a $\boldsymbol{q}_i$ vector for each objective o_i, thus also a vector of optimal actions $\boldsymbol{a}$. To obtain a single action, we may sum-up vectors $\boldsymbol{q}_i$ and then select the action corresponding to the maximum summed q-value:

$$a = a_j, \text{ for such } j \text{ that } \boldsymbol{z}_j = \max \boldsymbol{z}, \text{ where } \boldsymbol{z} = \sum_{i=0}^{N} \boldsymbol{q_i} \tag{5}$$

In this approach, q-values may be interpreted as votes of certain DQN, which are summed-up and the action with the highest score is selected. We need to stress here that simply adding the vectors does not produce a meaningful result yet. The q-values produced by different Q-functions are not scaled. In general q-values may be any real numbers. If we want them to represent votes for particular actions, each $\boldsymbol{q}_i$ vector needs to be rescaled to $[0,1] \subseteq \mathbb{R}$, where the $\min(\boldsymbol{q}_i)$ is mapped to 0 and $\max(\boldsymbol{q}_i)$ to 1.

Now, using the rescaled $\boldsymbol{q}_i$ vectors we can sum them up and select one action with the highest total q-value. For example, let have actions a_1, a_2, a_3, objectives o_1, o_2 and corresponding q-vectors $\boldsymbol{q}_1 = [0, 0.6, 1]$ and $\boldsymbol{q}_2 = [1, 0.5, 0]$. Adding them will result in vector $[1, 1.1, 1]$, for which the second element is the maximal, thus the corresponding action a_2 should be selected.

3.2 Suppression

We now have a method of combining outputs from several DQNs. Still, however such a combination will hardly lead to a meaningful action selection. Let us return to the previous example and say that we are driving a car; actions a_1, a_2, a_3 correspond to turning left, going straight, and turning right respectively. If we approach a wall and perform the action proposed in q_1 we will turn right, if we agree with the proposition in q_2 then will will turn left. Using the sum will

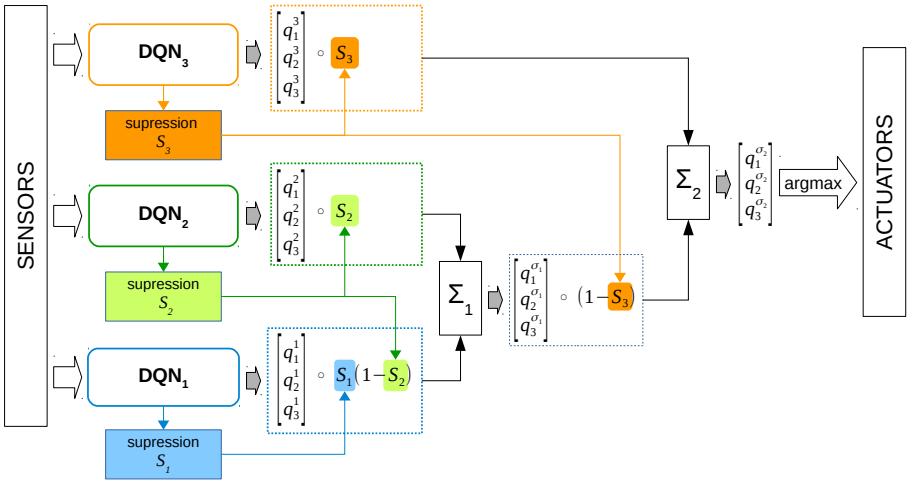

Figure 1. Three Deep Q-Networks are working in parallel based on the same sensory input. Each DQN corresponds to different task pursued by the agent. DQNs are stacked in a hierarchy, where DQN_1 has the lowest priority and DQN_3 the highest. Each DQN has an additional output which controls the suppression signal. The output of DQN_1 may be suppressed by DQN_2, while the summed output of DQN_1 and DQN_2 may be suppressed by DQN_3

however lead to going straight forward and hitting the wall. So while both DQNs suggested a meaningful action, their sum is not meaningful at all.

To solve this issue, we may use priorities assigned to objectives as described in Section 2. As each DQN in our model is corresponding to a particular objective, we may say that if DQN_i is providing q-values with respect to objective o_i, then this DQN has the same priority i as objective o_i. DQNs with priorities may be stacked to form a hierarchy. Now, we may sum two q-vectors with weights, such that the q-vector from DQN with higher priority will have a greater weight, and as a result - more impact on the action being selected. To control the value o this weights, we will introduce *suppression connections* - an idea borrowed from the subsumption architecture.

In our architecture, we enable DQNs with higher priority to *suppress* outputs of DQNs with lower priority. An example architecture of three DQNs stacked in a hierarchy with suppression connections is presented in Figure 1.

The suppression signal $S \in [0,1] \subseteq \mathbb{R}$ is a value that is used as a weight for summing the q-vectors. For example, if $\boldsymbol{q}_a$ is the output q-vector of DQN_a to be suppressed, $\boldsymbol{q}_b$ is the output q-vector of DQN_b that is suppressing (priority of DQN_b is higher than DQN_a) and S_b is the suppression signal, then the

suppressed sum of outputs is defined as:

$$\boldsymbol{q}_{\sigma_1} = (1 - S_b)\boldsymbol{q}_a + S_b\boldsymbol{q}_b \tag{6}$$

The interpretation of the suppression signal is straightforward. While S_b approaches 1 the output vector depends mostly on the values of $\boldsymbol{q}_b$ and in that case the output of DQN_a is suppressed. The opposite occurs while S_b approaches 0 and then the values of $\boldsymbol{q}_a$ are dominating in the output.

We can repeat the same procedure for an additional DQN. Let $\boldsymbol{q}_c$ be the output of DQN_c, which has suppression signal S_c and has priority higher than DQN_b and DQN_a. Then the output of the whole hierarchy will be:

$$\begin{aligned} \boldsymbol{q}_{\sigma_2} &= (1 - S_c)\boldsymbol{q}_{\sigma_1} + S_c\boldsymbol{q}_c \\ &= (1 - S_c)(1 - S_b)\boldsymbol{q}_a + (1 - S_c)S_b\boldsymbol{q}_b + S_c\boldsymbol{q}_c \end{aligned} \tag{7}$$

In a general case, we may have a N-level hierarchy of DQNs, such that a DQN at level n is suppressing the summed outputs of DQNs from levels $(0, \ldots, n-1)$. In such case the output of a N-level DQN hierarchy is equal to:

$$\boldsymbol{q}_\sigma = S_N\boldsymbol{q}_N + \sum_{i=1}^{N-1} S_i \left(\prod_{k=i+1}^{N} (1 - S_k) \right) \boldsymbol{q}_i \tag{8}$$

Additionally, the output q-vector may be summed with a random q-vector when the values of all suppression outputs are close to zero. The weight of the random vector is $\prod_{i=0}^{N}(1 - S_i)$. This ensures that a random action is selected when none of DQNs provide a substantial share to the output q-vector.

The suppression signals are provided by dedicated suppression outputs by each DQN respectively. In this work we restricted the model to one suppression output per DQN, however one may expand this idea to enable multiple suppression outputs from a single DQN.

3.3 Objectives, goals and rewards

A multi-objective agent receives rewards with respect to each objective. Objectives of an agent may derived from some envisioned behaviour of the agent. Here we will consider how a complex behaviour of an agent may be decomposed into several objectives.

First we should consider how to define objectives of the agent. For example, we may want a mobile robot to behave in such way to recharge its batteries automatically and move without colliding with other objects in the environment. The robot thus may have two objectives: a) to seek for energy sources and b) to avoid colliding with obstacles.

In the default situation, the robot consumes energy for moving and performing other tasks. The state when the robot is recharging is attractive for the robot, as being in this state enables the robot to realize objective (a). On the contrary, colliding with an obstacle is a state that the robot should avoid, as

being in this state clearly disallows the robot to realize objective (b). In case of embodied agents, especially robots, one can usually describe an objective as *attractive* or *repulsive*. In terms of rewards, an agent moving towards an attractive state should be rewarded, whereas an agent moving towards a repulsive state should be punished (receive a negative reward).

We can define types of objectives in more details as follows:

1. attractive - the goal is to **be in** a state or set of states, a positive reward is generated when a desired state is achieved, negative or zero reward is generated for all other states
2. repulsive - the goal is to **avoid** a state or set of states, a negative reward is generated when an undesired state is achieved, positive or zero reward is generated for all other states

Both the desired states in the attractive objectives and undesired states the repulsive objectives will be referred to as *goal states*. For each objective o_i, we may define a set of rewards $R = \{r_g, r_d\}$, where r_g is the reward for being in a goal state and r_d is the default reward for not being in the goal state. Note that r_g will be positive for attractive and negative for repulsive objectives respectively. Moreover we will define *goal* function as follows:

$$goal(r) = \begin{cases} 1 & \text{if } r = r_g \\ 0 & \text{if } r \neq r_g \end{cases} \tag{9}$$

The *goal* function will tell us whether received reward is related with achieving a goal state by the agent.

3.4 Learning

In our solution, we use Deep Q-Networks to approximate the values of Q-functions. Following the state of the art in this field a DQN provides the approximated function $Q(s, a; \theta)$, where θ are the learnable parameters of the neural network. As in our model we use multiple DQNs, there is a function $Q_i(s, a; \theta_i)$ for a DQN_i related to objective o_i. Each DQN_i is optimised iteratively, using the following loss function for each iteration j:

$$L^Q_{i,j}(\theta_{i,j}) = \mathbb{E}_{(s,a,r_i,s')\sim U(D_i)} \left[\left(r_i + \gamma \max_{a'} Q_i(s', a'; \theta^-_{i,j}) - Q_i(s, a; \theta_{i,j}) \right)^2 \right] \tag{10}$$

As introduced in [1], there are in fact two neural networks involved in the learning process of a single DQN. The *on-line network* $Q_i(s, a; \theta)$ is updated at each iteration, while the *target network* $Q_i(s', a'; \theta^-)$ is updated only each K iterations. Moreover *experience replay* is used to further improve the learning process. The agent stores experienced states, actions and rewards in a *replay memory* D_i for each DQN_i respectively. Then at each iteration, each DQN_i is trained using a sample of past experiences selected uniformly at random from the corresponding replay memory D_i. Those samples are used as mini-batches for gradient descent

optimization. There are more recent improvements to Deep Q-learning such as double q-learning or duelling networks, which were not used in our model, as a *vanilla* DQN was more plausible for evaluation and comparison.

In our model DQNs provide not only Q-values for actions, but also the values for the suppression signals S. Each DQN_i provides a suppression value $S_i(s;\theta)$, which depends on the state s. Let us here briefly analyse what are the desired values of the suppression signal. First, we want the suppression output to be active (to suppress lower-priority DQNs) when the agent should perform action moving him towards desired state with respect to an objective o_i. Informally speaking, the suppression signal of DQN_i should tell "how good" it is to use DQN_i for selecting the next action in state s. Let us define the suppression reward as: $\rho_i = goal(r_i)$. According to our definition of the *goal* function, such reward would return 1 only if the current state is the goal state; in all other cases it would return 0. We may perceive the suppression value as if it was a *state-value* function [3], returning the value of the state s under policy π:

$$S_\pi(s) = \mathbb{E}_\pi \left[\sum_{k=0}^{\infty} \gamma^k \, \rho_{i,t+k+1} \,\middle|\, s_t = s \right] \tag{11}$$

The policy π is the policy of the whole agent (a result of combining multiple Q-functions). The proposed S-value function will hence provide values representing the chances of achieving a goal state (with respect to some objective o) given the current state s and following policy π.

The suppression signal is used not only to select the output of particular DQN, but more importantly, also to disable outputs of DQNs with lower priority. However, a disabled DQN is not able to pursue its objective, thus we want to avoid situation when a DQN is suppressed to often and unnecessarily. To ensure this, we may simply punish the suppression output each time the suppression signal leads to a negative reward of a lower-priority DQN by extending the ρ_i definition:

$$\rho_i = goal(r_i) + \min(\sum_{k=1}^{i-1} S_i \cdot r_k, 0) \tag{12}$$

we will refer to this extended definition of ρ as *cross-reward*. Enabling cross-reward may improve the stability of learning, however it makes suppression signal dependent on a particular order of priorities.

The S-value may be updated using TD-learning. As we use a neural network for approximating $S_i(s)$, we may define the loss function as follows:

$$L^S_{i,j}(\theta_{i,j}) = \mathbb{E}_{(s,\rho_i,s')\sim U(D_i)} \left[\left(\rho_i + \gamma S_i(s';\theta_i,j^-) - S_i(s;\theta_{i,j}) \right)^2 \right] \tag{13}$$

The suppression value is provided by an additional output of the DQN and the learning procedure is analogical to Q-function. The neural network is optim-

ised using a combined loss function for both Q-values and S-values:

$$L_{i,j}(\theta_{i,j}) = \mathbb{E}_{(s,a,r_i,\rho_i,s')\sim U(D_i)}\Bigg[\Big(r_i + \gamma \max_{a'} Q_i(s', a'; \theta^-_{i,j}) - Q_i(s, a; \theta_{i,j})\Big)^2 + \big(\rho_i + \gamma S_i(s'; \theta^-_{i,j}) - S_i(s; \theta_{i,j})\big)^2\Bigg] \tag{14}$$

4 Evaluation

4.1 Eater - a 2D game-like virtual environment

To evaluate the solution presented in this paper we created *Eater* - a simple game-like virtual environment, which consists of an agent, walls and objects consumable by the agent. Eater is presented in Figure 2. The agent is a circular object that may move around the map by performing one of three actions: forward, forward and left turn, forward and right turn. The map is a continuous space. The agent perceives the environment only by visual sense, i.e. a W x H pixel (width and height) rectangle situated in front of him. This visual input is converted to gray-scale (8bit). Agent's world (white) is surrounded by walls (black), which agent can not pass. Agent may pick up two objects: *food* and *poison*. Food is indicated by three small coaxial circles (black), while poison indicated by a filled circle (grey). Food and poison re-spawn at random positions on the map. The quantity of each is constant during the game. Eater is a simplified simulation of a mobile robot moving on a flat surface (e.g. floor) with a video camera attached at the top of the robot pointed towards the floor.

The agent has an integer energy level $E < E_{max}$, which is decreased at each time step by E_{step}, until it reaches 0 and the game restarts (the agent is respawned at random position). The agent may collect food, which increases the energy level by E_{food}, or may collect poison which decreases the energy level by E_{poison}. Colliding with a wall decreases the agent's energy level by $E_{collision}$. The agent starts each game with initial $E = E_{start}$.

Depending on the variant of the game, the agent may have up to three objectives: a) avoid colliding with walls, b) gather food and c) avoid gathering poison. We will refer to game variants as: G2 - objectives (a) and (b), G3 - objectives (a),(b) and (c). One should note, that there is another indirect objective of the agent - maintain energy level above 0. For variants G2 and G3, we may say that the agent plays optimally, if the game never ends, thus the agent is able to maintain the $E > 0$ continuously.

The rewards for particular objectives are as follows: objective (a): -1 for collision, $+0.1$ otherwise; objective (b): $+1$ for gathering food, -0.1 otherwise; objective (c): -1 for gathering poison, $+0.1$ otherwise.

In all experiments described in this chapter, the game options were as follows: $E_{start} = 2000$, $E_{step} = 2$, $E_{collision} = 10$, $E_{food} = E_{poison} = 300$, $E_{max} = 5000$.

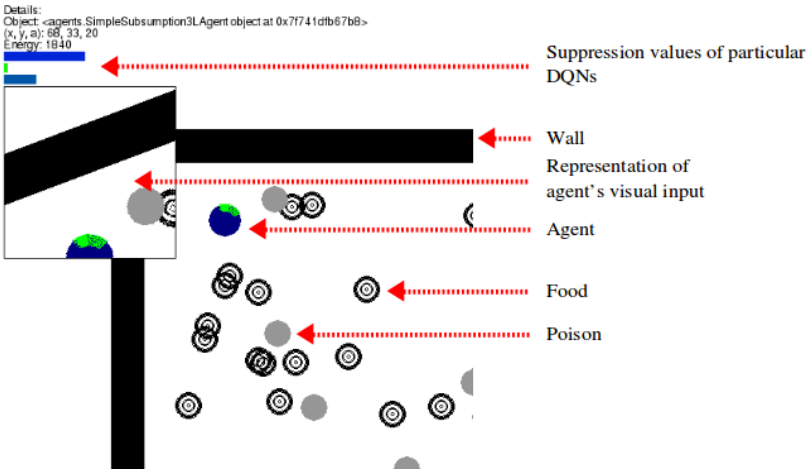

Figure 2. Eater - a game-like virtual environment with agent pursuing multiple-objectives. The environment consists of the agent, walls and consumables: food and poison. The agent perceives the environment by a visual input (a view from the top limited to a square located in the front of the agent). Agent may move forward and turn; its area of movement is limited by walls. Agent may have up to three objectives: avoid walls, consume food and avoid consuming poison.

The size of the game area is 500x500 pixels. The size of the agent sight rectangle is $W = 180\,\text{px}$, $H = 180\,\text{px}$. The quantity of food is 20 and poison is 8. At each step n the energy level of the agent is equal to:

$$\begin{aligned} E_0 &= E_{start} \\ E_{n+1} &= \min(E_n - E_{step} - (E_{collision}|\text{collision}) + \\ &\quad + (E_{food}|\text{food}) - (E_{poison}|\text{poison}), E_{max}) \end{aligned} \tag{15}$$

The agent performance is scored by two measures. First is the number of games in a given period - N_{games}. Second measure is the total sum of extra energy gathered/lost by the agent in a given period:

$$S_{energy} = \sum_{t}^{N} E_{food_t} - E_{collision_t} - E_{poison_t} \tag{16}$$

For a perfectly playing agent, N_{games} should be 0, as the agent would never be restarted. Higher values of S_{energy} correspond to better efficiency of the agent in avoiding or seeking goal states.

Note that the probabilities of agent reaching certain goal states are not equal. The probability of collision is much higher than the probability of gathering poison. Moreover, there is no linear relation between reward values and the gain/loss of the energy of the agent. This may seem to be an unnecessary complication, however we find this approach a good representation of real world

Table 1. mDQNS hyperparameters

Parameter	Value
replay memory size	100000
target network update rate	2000
learning rate	0.0025
ϵ start value	1
ϵ end value	0.1
ϵ end step	100000
discount	0.99
batch size	32
optimiser	RMSProp

problems found in e.g. robotics. The energy level of the agent may be easily mapped to the battery level of a mobile robot. A robot may be unaware of the precise value of energy loss (or some fitness score) experienced during a collision, but any such collision should be punished by the internal reward system of such robot, as in most cases it has a negative impact on the robot's performance.

4.2 mDQNS implementation

Our implementation of the mDQNS was based on a DQN implementation for TensorFlow[16] provided by [17]. We expanded the standard DQN with additional suppression outputs and mechanism for stacking several such expanded DQNs to form a mDQNS. Each single DQN in a mDQNS consist of a convolution network with three convolution layers and no pooling layers, followed by a fully connected layer and the output layer. The suppression output is a sigmoid layer connected to the fully connected layer. The parameters of the convolution network were based on the specification provided in [2]. Exception is the size of the fully connected layer, which in our experiments is 64, and the size of the input image which in our case is 60x60x4. Our implementation was not optimised for execution speed, however to make it less memory-consuming we introduced a common memory for storing past states used for memory replay (as all DQNs in mDQNS use the same visual input).

The hyperparameters used for training DQNs during evaluation are presented in Table 1. Some modifications were present: we used progressive values of the learning rate for DQNs on different layers in the hierarchy. The top DQN (with highest priority) used the learning rate value as presented in the table; for each next DQN lower in the hierarchy the learning rate was twice higher. The motivation is to speed up the behaviour learning by the lower layers, so higher layers may learn when to disable this behaviour.

4.3 Experiments

To evaluate our method we propose the following approach. First we measure performance in the G2 variant of Eater for: monolithic DQN (mono-DQN),

2-objective multiple DQN without subsumption (2-mDQNn) and 2-objective multiple DQN with subsumption (2-mDQNS), using hyperparameters resembling those in [2] and presented in Table 1.

Next, we switch to the G3 variant of Eater and measure mono-DQN, 3-objective multiple DQN (3-mDQNn) and 3-objective DQN with subsumption (3-mDQNS).

In case of 2-mDQNn and 3-mDQNn the outputs of particular DQNs are summed as described in Section 3.1. In those cases, suppression outputs of DQNs are not used and not trained. For 2-mDQNS and 3-mDQNS, the suppression outputs are used and combined as described in Section 3.2. The presented evaluation results are obtained from averaging 5 runs for each case. A video material showing the behaviour of the agent for particular evaluation cases is provided[1].

G2 variant - mono-DQN vs 2-mDQNn vs 2-mDQNS In this scenario we use G2 variant of the game. The agent has objectives (a) and (b) introduced in the previous section. The agent thus receives two rewards: r_1 and r_2 at each step. In case of the monolithic DQN, the agent receives a sum of rewards: $r_s = r_1 + r_2$. In the 2-mDQNS case, DQN_2 receives rewards for objective (b) - food collection and may suppress DQN_1, which receives rewards for objective (a) - collision avoidance.

In Figure 3 (left column), we may observe the number of games and score values during training for G2 case. The monolithic approach and 2-mDQNS achieve very similar results: the number of games drops to 0 after few thousand training steps, thus both methods learn to play optimally and they achieve almost identical score. The mDQNn, on the other hand, is not able to learn an optimal strategy and the agent is restarted frequently during training.

We may observe that, as expected, each DQN module in the mDQNS case learned a distinct behaviour. The suppression output of the DQN_2 responsible for food gathering gets activated (close to 1) when food object is in the field of view of the agent, in effect the agent performs actions suggested by DQN_2, which lead to food gathering. On the contrary, while approaching the wall, the suppression output of DQN_2 is inactive (close to 0) and actions suggested by DQN_1 drive the agent to avoid the wall.

G3 variant - mono-DQN vs 3-mDQNn vs 3-mDQNS In this variant the mDQNn and mDQNS consist of three DQNs: DQN_1, DQN_2, and DQN_3 receiving rewards for objectives (a), (b) and (c) respectively. Moreover, in case of mDQNS, DQN_2 suppresses DQN_1, while DQN_3 suppresses both DQN_2 and DQN_1. For the monolithic DQN, the rewards for particular objectives were summed up, thus $R = r_{o_1} + r_{o_2} + r_{o_3}$.

The experiment results are presented in Figure 3 (right column). In this case the monolithic approach gave the best results in terms of score and the number

[1] The video material attached to this paper will be uploaded to some video hosting service for the camera-ready version

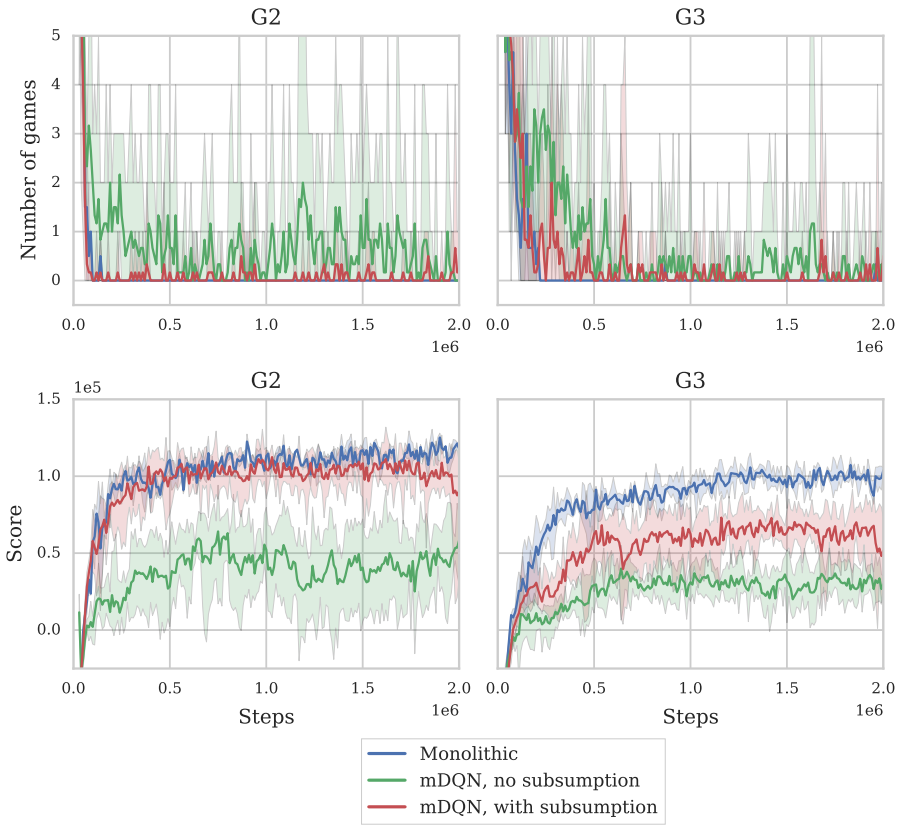

Figure 3. Evaluation results for two variants of the Eater game: G2 - two objectives, G3 - three objectives. The number of games and score is provided for: a) monolithic DQN, b) multiple DQN without subsumption (mDQNn) and c) multiple DQN with subsumption (mDQNS).

of games. The mDQNS solution required more steps to converge to near optimal solution and the achieved score was lower compared to the monolithic solution. Still, mDQNS achieved considerably better results than mDQNn, showing that suppression is indeed improving the learning in multiple DQNs.

Similarly as in G2 variant, agent driven by mDQNS was able to learn distinct behaviours. The added DQN_3 learned to avoid poison objects. The suppression output of DQN_3 became active when poison was located directly in front of the agent. The difference in score between mDQNS and the monolithic approach may be the result of a particular objective ordering. This aspect requires future evaluation.

In both G2 and G3 variants, it was possible to manipulate the behaviour of the mDQNS driven agent post-learning by forcing suppression outputs to certain

values. For example, agent in G3 variant, only avoids walls and poison, but does not gather food, when DQN_2 suppression output is forced to 0.

5 Conclusions

In this paper, we presented a method for using multiple Deep Q-Networks to approach multi-objective problems. We introduced suppression signals to DQNs in order to create a hierarchy of multiple DQNs, which is closely resembling the subsumption architecture. The implementation of Deep Q-Network was extended with the suppression output, which may be trained using rewards received by the agent. We refer to our solution as to Multiple Deep Q-Networks with Subsumption (mDQNS).

In the experimental part, we shown that mDQNS is able to successfully learn distinct behaviours related to particular objectives. The results of our experiments are promising and justify further development of our approach.

The presented solution has several advantages, useful in practical applications: it may learn multiple behaviours at once, based on high-level visual input; learned behaviours may be modified after learning by throttling the output of DQNs responsible for manifesting particular sub-behaviours; the agent may be easily extended with new behaviours by adding more layers to the hierarchy; for specific cases, the ordering of objectives may be changed at runtime.

In future work we want to improve the performance of mDQNS and evaluate it against various multi-objective problems or games. Moreover, mDQNS may be optimised by using common convolution layers shared by many DQNs. It is particularly interesting to use mDQNS in multi-agent environments, where it could learn behaviours such as e.g. predator-prey interaction. The implementation of mDQNS in a physical robot would be a challenge and could provide many useful insights.

Acknowledgements Author thanks Piotr Wasilewski for his scientific supervision and valuable suggestions. This work was supported by the National Science Centre in programme SONATA 1, grant no. 2011/01/D/ST6/06981.

References

1. Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., Riedmiller, M.: Playing atari with deep reinforcement learning. In: NIPS Deep Learning Workshop. (2013)
2. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., Hassabis, D.: Human-level control through deep reinforcement learning. Nature **518**(7540) (Feb 2015) 529–533 Letter.
3. Sutton, R.S., Barto, A.G.: Introduction to Reinforcement Learning. 1st edn. MIT Press, Cambridge, MA, USA (1998)

4. Liu, C., Xu, X., Hu, D.: Multiobjective reinforcement learning: A comprehensive overview. IEEE Transactions on Systems, Man, and Cybernetics: Systems **45**(3) (March 2015) 385–398
5. Moffaert, K.V., Drugan, M.M., Now, A.: Scalarized multi-objective reinforcement learning: Novel design techniques. In: 2013 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL). (April 2013) 191–199
6. Barrett, L., Narayanan, S.: Learning all optimal policies with multiple criteria. In: Proceedings of the 25th International Conference on Machine Learning. ICML '08, New York, NY, USA, ACM (2008) 41–47
7. Moffaert, K.V., Nowé, A.: Multi-objective reinforcement learning using sets of pareto dominating policies. Journal of Machine Learning Research **15** (2014) 3663–3692
8. Sprague, N., Ballard, D.: Multiple-goal reinforcement learning with modular sarsa(o). In: Proceedings of the 18th International Joint Conference on Artificial Intelligence. IJCAI'03, San Francisco, CA, USA, Morgan Kaufmann Publishers Inc. (2003) 1445–1447
9. Raicevic, P.: Parallel reinforcement learning using multiple reward signals. Neurocomputing **69**(1618) (2006) 2171 – 2179 Brain Inspired Cognitive SystemsSelected papers from the 1st International Conference on Brain Inspired Cognitive Systems (BICS 2004)1st International Conference on Brain Inspired Cognitive Systems (BICS 2004).
10. Brys, T., Harutyunyan, A., Vrancx, P., Taylor, M.E., Kudenko, D., Nowe, A.: Multi-objectivization of reinforcement learning problems by reward shaping. In: 2014 International Joint Conference on Neural Networks (IJCNN). (July 2014) 2315–2322
11. Mossalam, H., Assael, Y.M., Roijers, D.M., Whiteson, S.: Multi-objective deep reinforcement learning. CoRR **abs/1610.02707** (2016)
12. Pfeifer, R., Scheier, C.: Understanding Intelligence. MIT Press, Cambridge, MA, USA (2001)
13. Brooks, R.: A robust layered control system for a mobile robot. IEEE Journal on Robotics and Automation **2**(1) (Mar 1986) 14–23
14. Wicaksono, H.: Q learning behavior on autonomous navigation of physical robot. In: 2011 8th International Conference on Ubiquitous Robots and Ambient Intelligence (URAI). (Nov 2011) 50–54
15. Farahmand, A.M., Ahmadabadi, M.N., Araabi, B.N.: Behavior hierarchy learning in a behavior-based system using reinforcement learning. In: 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566). Volume 2. (Sept 2004) 2050–2055 vol.2
16. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: Large-scale machine learning on heterogeneous systems (2015) Software available from tensorflow.org.
17. Kim, T.: Tensorflow implementation of human-level control through deep reinforcement learning. `https://github.com/devsisters/DQN-tensorflow`